

The Hamiltonian Brain

Laurence Aitchison¹ and Máté Lengyel²

July 7, 2014

¹ Gatsby Computational Neuroscience Unit, UCL,

`laurence@gatsby.ucl.ac.uk`

² Computational & Biological Learning Lab, University of Cambridge,

`m.lengyel@eng.cam.ac.uk`

Abstract

A venerable history of models have shown that simple and complex cell responses in the primary visual cortex (V1) are adapted to the statistics of natural images. These models account for trial-averaged neural responses in V1, by assuming that responses represent maximum likelihood or *a posteriori* inferences in a generative model trained on natural images. However, such models have very different structural and dynamical properties from the brain. In particular, these models violate Dale's law, and have deterministic gradient ascent (or instantaneous) dynamics that decay to a single fixed point in response to fixed input, belying the dynamical richness of cortical responses involving both substantial inter- and intra-trial variability and oscillatory components. We give a solution to all of these problems, by suggesting that the brain samples possible explanations of incoming data using a mechanism inspired by Hamiltonian Monte Carlo. Our network has recurrent connections that obey Dale's law, gives rise to oscillatory dynamics, and continues stochastic sampling in response to fixed input, rather than reaching a single fixed point. These properties of our model allow it to match three key aspects of neural responses in V1 that otherwise seem mysterious from the perspective of probabilistic inference. First, the frequency of oscillations increases with stimulus contrast. Second, there are large transient increases in firing rate upon stimulus onset. Third, excitation and inhibition are balanced, and inhibition lags excitation. While previous work has suggested that cortical dynamics can be understood as performing approximate inference, and in particular as sampling a posterior distribution, our results are the first to point to the dynamical footprints of a specific class of sampling algorithms that the cortex might actually implement.

1 Introduction

V1 has long fascinated experimental and computational neuroscientists, and is regarded as one of the best understood cortical areas [1]. One particularly influential class of functional models for V1 define a probabilistic models of images and use gradient-ascent to find the most probable explanation of the presented image, i.e. perform maximum *a posteriori* inference [2]. However, these models suffer from both biological and computational problems. The brain is an oscillating network, composed of excitatory and inhibitory neurons [3], in contrast, these models use gradient-ascent (non-oscillatory) dynamics, and have neurons that violate Dale’s law [4] by having both excitatory and inhibitory outputs. Computationally, such models are limited because they find only one most probable explanation for the data, whereas in order to make decisions and integrate information from multiple sources, it is usually necessary to consider uncertainty [5]. In this paper, we take a first step towards an oscillating, excitatory-inhibitory (EI) model of V1, that takes uncertainty into account.

Perhaps the most fundamental problem with gradient ascent-based models is that, in response to a fixed stimulus, neural activity does not converge to a fixed point as gradient ascent models assume, but continues oscillating noisily. This observation suggests that the brain does not find only the most probable explanation for the data, but instead samples multiple plausible explanations. This alternative, known as the sampling hypothesis [6, 7], not only accounts for multiple types of variability, but also correctly predicts that stimulus onset quenches neural variability [8], and that spontaneous activity should reflect the statistical properties of evoked activity [9]. Furthermore, sampling provides the most general representation of uncertainty of any generic machine learning method, so it has become central to many machine learning applications [10].

While the sampling hypothesis makes strong predictions about the statistics of stationary activity (that is, snapshots of activity patterns once the system has reach its stationary distribution), it makes no predictions about neural dynamics, including transients. To make such predictions, we need to specify a sampler. Many neurally plausible samplers have been suggested for discrete variables [11–13], but only one has been proposed for continuous variables, Langevin sampling [14, 15]. However, Langevin samplers have noisy gradient-ascent dynamics, so they not only retain the problems of maximum likelihood models, but also sample slowly [16, 17].

The failure of previously applied sampling methods lead us to think about Hamiltonian Monte Carlo (HMC). In its most general form, HMC samples the distribution over latent variables, \mathbf{u} , given observations, \mathbf{x} , by introducing an auxiliary variable, \mathbf{v} , with some conditional distribution, $P(\mathbf{v}|\mathbf{u}, \mathbf{x})$. Samples from the joint distribution, $P(\mathbf{u}, \mathbf{v}|\mathbf{x})$, are computed by running the dynamical

system,

$$\dot{\mathbf{u}} = -\frac{1}{\tau} \frac{\partial \log P(\mathbf{u}, \mathbf{v}|\mathbf{x})}{\partial \mathbf{v}} \quad (1)$$

$$\dot{\mathbf{v}} = -\frac{1}{\tau} \frac{\partial \log P(\mathbf{u}, \mathbf{v}|\mathbf{x})}{\partial \mathbf{u}}. \quad (2)$$

However, these dynamics are not ergodic as they hold $\log P(\mathbf{u}, \mathbf{v}|\mathbf{x})$ constant; to ensure ergodicity, noise must be added to either \mathbf{u} or \mathbf{v} . The standard machine learning formulation uses $P(\mathbf{v}|\mathbf{u}, \mathbf{x}) = P(\mathbf{v}) = \mathcal{N}(\mathbf{v}; \mathbf{0}, \mathbf{M})$, and adds noise by resampling \mathbf{v} .

We begin Section 2 by describing our HMC-inspired sampler. In order to be able to make specific predictions, we needed to choose a concrete probabilistic model. A Gaussian scale mixture (GSM) model was convenient, as, despite its simplicity, it captures some of the more interesting statistical properties of images [18], and gives a good description of the trial-averaged responses of simple cells [19]. We demonstrated that our dynamics sample the GSM model more rapidly than Langevin dynamics. Next we showed, by both theoretical analysis and simulation, that our sampler reproduces three properties of observed neural dynamics. First, our sampler has balanced excitation and inhibition, with inhibition lagging excitation [20]. Second, our sampler oscillates, and the oscillation frequency increases with stimulus contrast [21, 22]. Finally there is a large, transient increase in firing rates upon stimulus onset [21]. Importantly, all these phenomena emerge naturally from the same normatively justified dynamics without having to resort to additional heuristics.

2 Results

2.1 The sampler

To make biologically motivated choices for the distribution of the latent variable, $P(\mathbf{v}|\mathbf{u}, \mathbf{x})$, and the noise process, we need to suggest how the brain represents \mathbf{u} and \mathbf{v} . For most probabilistic models, elements of \mathbf{u} and \mathbf{v} are dimensionless, $O(1)$ quantities with no lower bound. We therefore represent \mathbf{u} and \mathbf{v} as membrane potentials, $\sigma\mathbf{u}$ and $\sigma\mathbf{v}$, where σ is simply a scaling factor with units of mV. Moreover, we expect \mathbf{u} , the variables of interest, to be represented by excitatory cells, as these cells have the potential to form long-range projections to other brain areas [3], and \mathbf{v} , the auxiliary variables, to be represented by local inhibitory interneurons.

To derive the dynamics, we needed to choose the prior, $P(\mathbf{u})$, the distribution of the auxiliary variable, $P(\mathbf{v}|\mathbf{u}, \mathbf{x})$, and the noise process. For simplicity, we took the prior over \mathbf{u} to be Gaussian, $P(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{C})$. Remembering that excitation and inhibition are balanced [20], we see that \mathbf{v} should, in some sense, be close to \mathbf{u} . Formally, we could encode this intuition by setting $P(\mathbf{v}|\mathbf{u}, \mathbf{x}) =$

$P(\mathbf{v}|\mathbf{u}) = \mathcal{N}(\mathbf{v}; \mathbf{u}, \mathbf{M})$, which differs from the usual choice used in machine learning both by having non-zero mean, and by depending on \mathbf{u} ; it is this choice that gives rise to many of the interesting properties of our sampler. For simplicity and tractability, we chose to add uncorrelated Gaussian white noise to both \mathbf{u} and \mathbf{v} . These choices yield the following dynamics:

$$\dot{\mathbf{u}} = \mathbf{W}_{\mathbf{uu}}\mathbf{u} - \mathbf{W}_{\mathbf{uv}}\mathbf{v} + \frac{1}{\tau_L} \frac{\partial \log P(\mathbf{x}|\mathbf{u})}{\partial \mathbf{u}} + \sqrt{\frac{2}{\tau_L}} \boldsymbol{\eta}_{\mathbf{u}} \quad (3a)$$

$$\dot{\mathbf{v}} = \mathbf{W}_{\mathbf{vu}}\mathbf{u} - \mathbf{W}_{\mathbf{vv}}\mathbf{v} - \frac{1}{\tau_H} \frac{\partial \log P(\mathbf{x}|\mathbf{u})}{\partial \mathbf{u}} + \sqrt{\frac{2}{\tau_L}} \boldsymbol{\eta}_{\mathbf{v}}, \quad (3b)$$

where,

$$\mathbf{W}_{\mathbf{uu}} = \left(\frac{1}{\tau_H} - \frac{1}{\tau_L} \right) \mathbf{M}^{-1} - \frac{1}{\tau_L} \mathbf{C}^{-1} \quad (3c)$$

$$\mathbf{W}_{\mathbf{uv}} = \left(\frac{1}{\tau_H} - \frac{1}{\tau_L} \right) \mathbf{M}^{-1} \quad (3d)$$

$$\mathbf{W}_{\mathbf{vu}} = \left(\frac{1}{\tau_H} + \frac{1}{\tau_L} \right) \mathbf{M}^{-1} \quad (3e)$$

$$\mathbf{W}_{\mathbf{vv}} = \left(\frac{1}{\tau_H} + \frac{1}{\tau_L} \right) \mathbf{M}^{-1} + \frac{1}{\tau_H} \mathbf{C}^{-1} \quad (3f)$$

We take $\frac{\partial \log P(\mathbf{x}|\mathbf{u})}{\partial \mathbf{u}}$ to be the external input to the network, so for the recurrent dynamics to obey Dale's law, the elements of the weight matrices (e.g. $\mathbf{W}_{\mathbf{uv}}$) must be positive. As long as $\tau_L > \tau_H$, this condition is easily achievable, as \mathbf{M}^{-1} is an arbitrary positive definite matrix, so its elements can be set to be as large and positive as necessary in order to cancel any negative contribution from \mathbf{C} . In our simulations, we set $\mathbf{C} = \mathbf{I}$ and $\mathbf{M} = \sigma_{\mathbf{v}}^2 \mathbf{I}$, where $\sigma_{\mathbf{v}}^2$ was small enough for Dale's law to be obeyed.

The parameters, τ_L and τ_H denote the timescales of the noisy, Langevin dynamics and the deterministic, Hamiltonian dynamics respectively. The third parameter, $\sigma_{\mathbf{v}}^2$, controls not only the degree of balance, but also the timescale of the recurrent connections, with smaller $\sigma_{\mathbf{v}}^2$ leading to faster timescales and tighter balance. To choose values for τ_L , τ_H and $\sigma_{\mathbf{v}}^2$, we considered biological constraints.

First, we estimated τ_L by comparing that the rate at which membrane potential variance increases in our equations, $2\sigma^2/\tau_L$, to the rate of increases given by stochastic vesicle release, the primary source of 'noise' in cortical circuits. If a neuron is connected to s presynaptic neurons, firing with average rate r , and the variance of a unitary EPSP is v , then stochastic vesicle release introduces variance at the rate srv . Setting $srv = 2\sigma^2/\tau_L$ allows us to find the Langevin timescale,

$$\tau_L = \frac{2\sigma^2}{srv}. \quad (4)$$

However, estimating τ_L is difficult, because there are huge uncertainties in σ , s , r and v . We therefore wrote our uncertainty about each parameter as a log-normal distribution, $P(\log x) = \mathcal{N}(\log x; \mu_x, \sigma_x^2)$ where x is one of σ , s , r , or v , and computed the induced distribution on τ_L . To specify the distributions, we wrote a range, from x_l to x_h , that, we believe, contains around 95% of the probability mass, taking the boundaries of the range to be two standard-deviations from the mean in the log-domain, $\log x_l = \mu_x - 2\sigma_x$ and $\log x_h = \mu_x + 2\sigma_x$.

To estimate the required ranges, we took values from the neuroscience literature. First, estimates of firing rates vary widely, from around 0.5 Hz [23] to around 10 Hz [24]. Second, the number of synapses per cell is usually taken to be around 10000. However, it is likely that there are multiple synapses per connection [25], so there could be anywhere from 1000 to 10000 input cells for a single downstream neuron. Third, the average variance per spike is relatively easy to measure, data from Song *et al.* [26] put the value at 0.076 mV^2 . As other measurements seem roughly consistent [27], we use a relatively narrow range for v , from 0.05 mV^2 to 0.1 mV^2 . Finally, the scaling factor, σ , could plausibly range from 2.5 mV to 7.5 mV, giving a full (2 standard deviations, and both sides of the mean) range of membrane potential fluctuations of 10 mV to 30 mV [28].

These ranges give a central estimate of $\tau_L = 150 \text{ ms}$, which we used in our simulations. In agreement with this back-of-the-envelope calculation, we find that our sampler’s dynamics match neural dynamics when τ_L lies in a broad range, from around 60 ms to around 400 ms (see Figure S1). While τ_L appears relatively large in comparison with typical neural timescales, which are often around 10 ms, it should be remembered that τ_L parameterises only the amount of noise injected into the network at every time step. The Langevin timescale does not therefore have any necessary link to other neural time constants, and could potentially be much larger than our estimate of 150 ms.

Finally, we set the remaining parameters, τ_H and σ_v^2 by noting that they parametrise not only timescales of abstract Langevin and Hamiltonian dynamical processes, but also the timescales of recurrent and external input to our network. The timescale of the external input term in Equation (3b) is τ_H , and the timescale associated with the average recurrent weights (ignoring the contribution from \mathbf{C}) is $\sigma_v^2 \tau_H$. In order to ensure that the dynamics remain biologically realistic, we constrain these two timescales to be above 10 ms. With all else being equal, shorter timescales lead to faster dynamics, so we set the two timescales mentioned above to their minimum values, which gave $\tau_H = 10 \text{ ms}$, and $\sigma_v^2 = 1$.

2.2 Gaussian scale mixture model

To simulate the dynamics we define a probabilistic model from which we can actually draw samples. We chose a Gaussian scale mixture model as, despite its simplicity, it still captures some of the more interesting statistical properties of

sensory natural image patches [18]. Our GSM model has two latent variables, \mathbf{u} and z , where $P(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{C})$ and $P(z) = \mathcal{N}(z; 0, 1)$. The observations, \mathbf{x} , are generated by,

$$P(\mathbf{x}|\mathbf{u}, z) = \mathcal{N}(\mathbf{x}; |z| \mathbf{A}\mathbf{u}, \sigma_{\mathbf{x}}^2 \mathbf{I}). \quad (5)$$

When the GSM is used to model image patches, \mathbf{x} represents the intensity of each pixel, \mathbf{u} represents the edges, or features that are being used in the image, and $|z|$ represents the average contrast level in the patch.

Not only does \mathbf{u} display similar mean responses to V1 simple cells [19], the GSM also has interesting implications for neural variability. To look at these implications, we would ideally like to write down $P(\mathbf{u}|\mathbf{x})$, but as $P(\mathbf{u}|\mathbf{x})$ cannot be written analytically, we instead look at $P(\mathbf{u}|z, \mathbf{x})$,

$$P(\mathbf{u}|z, \mathbf{x}) = \mathcal{N}\left(\mathbf{u}; \frac{z}{\sigma_{\mathbf{x}}^2} \Sigma(z) \mathbf{A}^T \mathbf{x}, \Sigma(z)\right), \quad (6)$$

where $\Sigma(z) = \left(\mathbf{C}^{-1} + \frac{z^2}{\sigma_{\mathbf{x}}^2} \mathbf{A}^T \mathbf{A}\right)^{-1}$, or in one-dimension,

$$\Sigma(z) = \left(C^{-1} + \frac{z^2}{\sigma_{\mathbf{x}}^2} A^2\right)^{-1}. \quad (7)$$

As the contrast level used to generate the image patch, z_{gen} , rises, the inferred contrast, z , rises, causing the posterior variance, Σ , to fall. If the sampling hypothesis is correct, then lower posterior variance should give rise to lower observed neural variability. The sampling hypothesis therefore provides an explanation for the observation that neural variability falls in response to stimulus onset [29] (at stimulus onset, the effective contrast rises from 0 to some non-zero value). This fundamental agreement between the probabilistic model and patterns of real neural activity is another key motivation for choosing to use the GSM model in this paper.

For our theoretical analyses of dynamics, it will prove useful to reduce the n -dimensional GSM to a one-dimensional toy model, as we did in Equation (7). To make the reduction, it is useful to have \mathbf{x} and \mathbf{u} being the same length. However, if \mathbf{x} is a large image, or if \mathbf{u} forms a highly overcomplete representation, then the number of dimensions in \mathbf{u} and \mathbf{x} may be very different. In order to aid interpretability, can think of \mathbf{x} in this case we as representing (eg. Gabor) filter activations rather than pixels (where filter activations and pixels are linearly related under the generative model). We show in the supplementary information that models of images and filter activations are equivalent, in the sense that for any GSM model of images, a GSM model of filter activations can be constructed that gives the same posterior for any image.

Finally, we chose parameter settings for the GSM model. We set $\mathbf{C} = \mathbf{I}$, $\sigma_{\mathbf{x}}^2 = 0.1$, and $n = 15$, where n is the dimensionality of \mathbf{u} , \mathbf{v} . While we could have

used $\mathbf{A} = \mathbf{I}$ to represent orthogonal filters, in reality, filters are not usually exactly orthogonal, which we represent by adding noise to the elements of \mathbf{A} ,

$$P(A_{ij}) = \mathcal{N}(A_{ij}; \delta_{ij}, 0.05^2). \quad (8)$$

2.3 Simulations

We simulated stimulus onset by first running the sampler until it reached equilibrium with “no stimulus”, then we ran the sampler for a fixed period of time in the “stimulus” condition. To represent “no stimulus” we generated the stimulus \mathbf{x} by sampling from $P(\mathbf{x}|z=0)$, and to represent “stimulus”, we sampled \mathbf{x} from $P(\mathbf{x}|z=z_{\text{gen}})$, where $z_{\text{gen}} \in \{0.5, 1, 2\}$ for stimuli with increasing contrast levels. The brain does not know z , so it must infer z together with \mathbf{u} , we therefore inferred z and \mathbf{u} in parallel, using our HMC inspired sampler.

To make contact with experimental data, we also needed to compute local field potentials (LFPs), and firing rates. There are many methods for computing LFPs, we chose the simplest, averaging the membrane potentials across neurons, as it gave similar results to the other methods we tried, and has no tunable parameters. To compute firing rates, we used a rectified linear function of the membrane potential:

$$f_i(t) = \begin{cases} u_i(t) & \text{if } u_i(t) > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

2.4 Comparing samplers

It is now possible to compare our sampler with a Langevin sampler. The fundamental constraint is the noise level, estimated above, of $\tau_L = 150$ ms. We found that the mean squared error between the true mean and the mean estimated using samples taken over 100 ms was around 0.3 for our Hamiltonian inspired sampler, and around 1.9 for a Langevin sampler. This difference indicates that our HMC inspired sampler uses limited noise far more efficiently than Langevin dynamics. This additional efficiency leads our sampler (Figure 1A) to move across the range of plausible values for $u_1(t)$ more frequently than the Langevin sampler (Figure 1B).

2.5 Dynamics

Our sampler’s dynamics match three important aspects of neural dynamics, inhibition balances excitation, there are oscillations with a stimulus dependent frequency, and a large transient occurs upon stimulus onset.




Figure 1: The Hamiltonian sampler is more effective than a Langevin sampler. Membrane potential traces for the Hamiltonian sampler (**A**) and the Langevin sampler (**B**).




Figure 2: Excitation and inhibition are balanced. In this plot, we used standard parameter settings, but with $n = 100$, to reduce noise by averaging over more cells. **A.** Average excitatory membrane potentials, $\langle u(t) \rangle$, and inhibitory membrane potentials, $\langle v(t) \rangle$, are positively correlated. **B.** Changes in average inhibitory membrane potentials lag behind changes in average excitatory membrane potentials.

2.5.1 Balance

Our model was premised on the observation that inhibition and excitation are close, or balanced [20], a property that is visible in simulation (Figure 2AB). Interestingly, inhibition lags excitation (Figure 2B) — a commonly observed property of neural systems [30] that we did not build in to our model.

2.5.2 Oscillations

Oscillations are a ubiquitous property of cortical dynamics [31], but in the context of V1 there are also specific, stimulus-induced patterns in these oscillations: the frequency of gamma-band oscillations increases with stimulus contrast [21, 22]. To analyse oscillations in our sampler we simplified the dynamics (Equation (3)) by looking at a one dimensional system, by eliminating the noise

(setting $\tau_L = \infty$), and by setting $P(u|x) = \mathcal{N}\left(u; c, (C^{-1} + A^2 z^2 / \sigma_{\mathbf{x}}^2)^{-1}\right)$,

$$\dot{u} = \frac{1}{\sigma_v^2 \tau_H} (u - v) \quad (10)$$

$$\dot{v} = \frac{1}{\sigma_v^2 \tau_H} (u - v) + \frac{1}{\tau_H} (C^{-1} + A^2 z^2 / \sigma_{\mathbf{x}}^2) (u - c). \quad (11)$$

This dynamical system can be rearranged to give the equation for simple harmonic motion,

$$\ddot{u} = \frac{1}{\sigma_v^2 \tau_H} (\dot{u} - \dot{v}) = -\frac{1}{\sigma_v^2 \tau_H^2} (C^{-1} + A^2 z^2 / \sigma_{\mathbf{x}}^2) (u - c). \quad (12)$$

Our sampler therefore oscillates, and the oscillation frequency increases with inferred stimulus contrast, z ,

$$f(z) = \frac{\sqrt{C^{-1} + A^2 z^2 / \sigma_{\mathbf{x}}^2}}{2\pi \tau_H \sigma_v}. \quad (13)$$

The increase in oscillation frequency is seen in simulated membrane potentials (Figure 3A–C) and simulated LFPs (Figure 3D–G). Furthermore, the quantitative predictions made by Equation (13) are close to simulated oscillation frequencies (Figure 3H).

2.5.3 Transients

Our simulations (Figure 4A) show large, stimulus contrast-scaled, transient increases in average firing rate at stimulus onset. These transients are a widely observed characteristic of responses in V1 [21] (as well as other sensory cortices). Interestingly, these transients are not trivially predicted by simpler models. Langevin sampling, for instance, does not give rise to any transient increase in firing rate — rates simply rise or fall towards their new steady state (Figure 4B). Moreover, if we reduce our sampler’s dynamics to noisy simple harmonic motion, by assuming that the brain knows z_{gen} , and sets $z = z_{\text{gen}}$, we still get no transient (Figure 4C).

To understand why our sampler gives transients, we therefore have to understand the interaction between the dynamics of \mathbf{u} and the inferred value of z . The generative model tells us that $x \approx zAu$, so we can write $z \approx x / (Au)$. This can be written in terms of z_{gen} rather than x , by assuming that $u_{\text{gen}} \approx 1$, so $x \approx Az_{\text{gen}}$, giving $z \approx z_{\text{gen}} / u$. Substituting for z in the simplified dynamics (Equation (3)) gives,

$$\tau \ddot{u} = -\frac{1}{\sigma_v^2 \tau_H^2} (C^{-1} + A^2 z_{\text{gen}}^2 / (\sigma_x^2 u^2)) (u - c). \quad (14)$$

Simulating this simplified dynamical system did indeed give large transients (Figure 4D). The key difference between these dynamics and simple harmonic motion is that the stiffness, $k(u) = C^{-1} + A^2 z_{\text{gen}}^2 / (\sigma_x^2 u^2)$ is not fixed, but varies

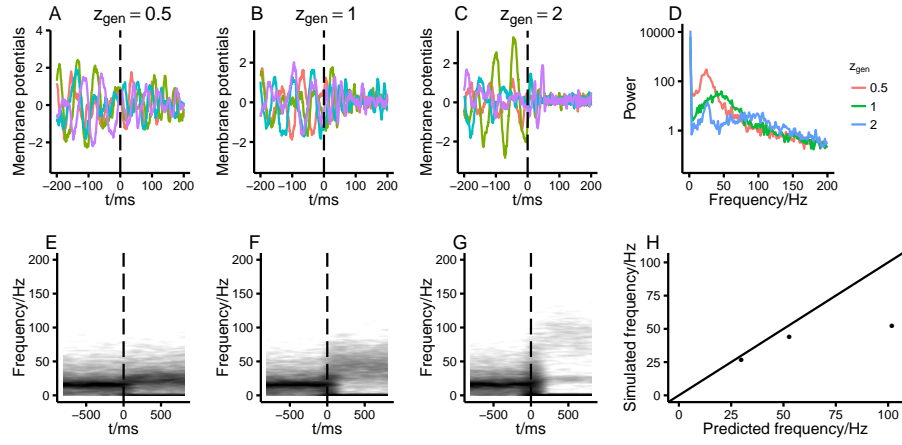


Figure 3: Oscillation frequency depends on stimulus contrast. **A, B, C.** The membrane potential response of one neuron to stimulus onset across 4 trials (coloured curves) shows that the variability decreases and the frequency increases as stimulus contrast increases (**A**, $z = 0.5$; **B**, $z = 1$; **C**, $z = 2$). **D.** Fourier transforms of the LFP, showing that oscillation frequency increases as z increases. **E–G.** Temporal Fourier transforms (Gaussian window, width 100 ms) of the LFP, with darker colours indicating higher power (contrast levels as in **A–C**). **H.** The simplified dynamics (x-axis, Equation (13)) predicts the qualitative trends of oscillation frequency in the simulations (y-axis).




Figure 4: Large firing rate transients in the model. **A.** Large transients are observed under the full dynamics. **B.** No transients are observed using a Langevin sampler ($\tau_L = 150$ ms). **C.** No transients are observed under the full dynamics when the inferred value of z is fixed, $z = z_{\text{gen}}$. **D.** Transients computed using the simplified dynamics (Equation (14)) are similar to transients under the full dynamics. Here, $c = A = C^{-1} = 1$, $\sigma_{\mathbf{x}}^2 = 0.1$, and the initial conditions are $u(0) = 0.1$ and $\dot{u}(0) = 0$. **E** The spring constant, with the same parameter settings as used in **D**.

with u . More precisely, the stiffness falls as the magnitude of u increases (Figure 4E). Therefore, just after stimulus onset, u is small, so there is a large force in the positive direction (due to the large stiffness), causing a large acceleration. Eventually, u exceeds c , but by that point the spring constant, and hence the restoring force has fallen, so the system’s momentum allows it to move a long distance, certainly further than if the spring constant were fixed. Our simplified model matches full simulations (Figure 4A) and the data [21] closely, both in terms of the transient timescale and the relative sizes of the transients for different contrast levels (values of z_{gen}).

3 Discussion

Despite being theoretically well-motivated, current probabilistic models of V1 are extremely unrealistic. They have recurrent dynamics that violate Dale’s law, display gradient-ascent like rather than oscillatory dynamics, and reach steady state in response to fixed input. In order to develop a more realistic model of neural dynamics, we took the observation that the brain continues to noisily sample new states in response to fixed input, and used it to suggest that the brain does not find one best explanation of incoming data, but samples multiple plausible explanations. While, this approach, known as the sampling hypothesis, makes interesting predictions about snapshots of activity patterns, making predictions about dynamics required us to specify a sampler. The only sampling method for continuous variables used so far in neuroscience, Langevin sampling, follows noisy gradient ascent dynamics, so it inherits the problems of gradient ascent models, and, in addition, samples slowly. We therefore developed a neurally plausible sampler inspired by HMC. Not only does the sampler oscillate, obey Dale’s law, and sample more rapidly than Langevin dynamics, it also matches three further aspects of neural dynamics, excitation and inhibition

are balanced, oscillation frequency increases with stimulus contrast [21], and large transients occur upon stimulus onset [21].

There are two important directions for further work.

First it is important to make a more detailed comparison of HMC dynamics to neural dynamics. We expect that as we improve the statistical model of images, HMC dynamics should reproduce an increasingly larger range of observed dynamical properties. For instance, in a complex statistical model that segments images into objects, we expect there to be strong statistical relationships between variables describing a single object. These strong statistical relationships should lead to strong dynamical relationships, including perhaps synchronisation, between variables associated with a single object [32].

Second, it is important to know how the brain might learn an HMC sampler, and might learn an effective, or fast, HMC sampler. There might, for instance, be two sets of learning rules running in parallel, one set of rules which learns that statistical structure of the input, and another that learns how to draw independent samples rapidly from the learned model. We would expect both of these sets of learning rules to involve inhibitory plasticity as an essential component [33].

In sum, we have given an approach by which successful, normative models of trial averaged activity in V1 can be extended to match the dynamics of neural activity, and to a lesser extent, the structure of neural connectivity.

Acknowledgements

This work was supported by the Wellcome Trust (ML), the Gatsby Charitable Foundation (LA), and the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement no. 269921 (BrainScaleS) (ML).

References

- [1] B. A. Olshausen and D. J. Field, “How close are we to understanding v1?,” *Neural Computation*, vol. 17, no. 8, p. 1665–1699, 2005.
- [2] B. A. Olshausen and D. J. Field, “Emergence of simple-cell receptive field properties by learning a sparse code for natural images,” *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.
- [3] E. R. Kandel, J. H. Schwartz, and T. M. Jessell, *Principles of neural science*, vol. 4. McGraw-Hill New York, 2000.
- [4] H. Dale, “Pharmacology and nerve-endings,” *Proceedings of the Royal Society of Medicine*, vol. 28, no. 3, pp. 319–332, 1935.

- [5] M. O. Ernst and M. S. Banks, “Humans integrate visual and haptic information in a statistically optimal fashion,” *Nature*, vol. 415, no. 6870, pp. 429–433, 2002.
- [6] P. O. Hoyer and A. Hyvarinen, “Interpreting neural response variability as monte carlo sampling of the posterior,” *Advances in neural information processing systems*, p. 293–300, 2003.
- [7] J. Fiser, P. Berkes, G. Orbán, and M. Lengyel, “Statistically optimal perception and learning: from behavior to neural representations,” *Trends in Cognitive Sciences*, vol. 14, no. 3, pp. 119–130, 2010.
- [8] G. Orbán, P.-O. Polack, P. Golshani, and M. Lengyel, “Stimulus-dependence of membrane potential and spike count variability in v1 of behaving mice,” *COSYNE Poster*, 2013.
- [9] P. Berkes, G. Orbán, M. Lengyel, and J. Fiser, “Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment,” *Science*, vol. 331, no. 6013, pp. 83–87, 2011.
- [10] C. Robert and G. Casella, “A short history of Markov Chain Monte Carlo: subjective recollections from incomplete data,” *Statistical Science*, vol. 26, no. 1, p. 102–115, 2011.
- [11] L. Buesing, J. Bill, B. Nessler, and W. Maass, “Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons,” *PLoS computational biology*, vol. 7, no. 11, 2011.
- [12] B. Nessler, M. Pfeiffer, L. Buesing, and W. Maass, “Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity,” *PLoS computational biology*, vol. 9, no. 4, 2013.
- [13] C. Savin, P. Dayan, and M. Lengyel, “Correlations strike back (again): the case of associative memory retrieval,” in *Advances in Neural Information Processing Systems*, p. 288–296, 2013.
- [14] R. Moreno-Bote, D. C. Knill, and A. Pouget, “Bayesian sampling in visual perception,” *Proceedings of the National Academy of Sciences*, p. 12491–12496, 2011.
- [15] A. Grabska-Barwinska, J. Beck, and P. E. Latham, “Demixing odors – fast inference in olfaction,” in *Advances in Neural Information Processing Systems*, 2013.
- [16] R. Neal, “MCMC for using hamiltonian dynamics,” *Handbook of Markov Chain Monte Carlo*, p. 113–162, 2011.
- [17] M. Lengyel, G. Hennequin, and L. Aitchison, “Fast sampling in recurrent neural circuits,” *COSYNE Poster*, 2014.
- [18] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, “Adaptive wiener denoising using a gaussian scale mixture model in the wavelet do-

- main,” in *International Conference on Image Processing*, vol. 2, p. 37–40, IEEE, 2001.
- [19] O. Schwartz and E. P. Simoncelli, “Natural signal statistics and sensory gain control,” *Nature Neuroscience*, vol. 4, no. 8, p. 819–825, 2001.
 - [20] M. Okun and I. Lampl, “Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities,” *Nature neuroscience*, vol. 11, no. 5, p. 535–537, 2008.
 - [21] S. Ray and J. H. Maunsell, “Differences in gamma frequencies across visual cortex restrict their possible use in computation,” *Neuron*, vol. 67, no. 5, pp. 885–896, 2010.
 - [22] M. J. Roberts, E. Lowet, N. M. Brunet, M. Ter Wal, P. Tiesinga, P. Fries, and P. De Weerd, “Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching,” *Neuron*, vol. 78, no. 3, p. 523–536, 2013.
 - [23] K. Mizuseki and G. Buzsáki, “Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex,” *Cell Reports*, vol. 4, no. 5, pp. 1010–1021, 2013.
 - [24] D. H. O’Connor, S. P. Peron, D. Huber, and K. Svoboda, “Neural activity in barrel cortex underlying vibrissa-based object localization in mice,” *Neuron*, vol. 67, no. 6, p. 1048–1061, 2010.
 - [25] T. Branco and K. Staras, “The probability of neurotransmitter release: variability and feedback control at single synapses,” *Nature Reviews Neuroscience*, vol. 10, no. 5, p. 373–383, 2009.
 - [26] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii, “Highly nonrandom features of synaptic connectivity in local cortical circuits,” *PLoS biology*, vol. 3, no. 3, p. e68, 2005.
 - [27] A. Bremaud, D. C. West, and A. M. Thomson, “Binomial parameters differ across neocortical layers and with different classes of connections in adult rat and cat neocortex,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 35, p. 14134–14139, 2007.
 - [28] E. A. Stern, A. E. Kincaid, and C. J. Wilson, “Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo,” *Journal of Neurophysiology*, vol. 77, no. 4, pp. 1697–1715, 1997.
 - [29] M. M. Churchland, B. M. Yu, J. P. Cunningham, L. P. Sugrue, M. R. Cohen, *et al.*, “Stimulus onset quenches neural variability: a widespread cortical phenomenon,” *Nature Neuroscience*, vol. 13, no. 3, pp. 369–378, 2010.
 - [30] E. O. Mann and I. Mody, “Control of hippocampal gamma oscillation frequency by tonic inhibition and excitation of interneurons,” *Nature Neuroscience*, vol. 13, no. 2, pp. 205–212, 2010.

- [31] G. Buzsaki, *Rhythms of the Brain*. Oxford University Press, 2006.
- [32] W. Singer, “Neuronal synchrony: a versatile code for the definition of relations?,” *Neuron*, vol. 24, no. 1, p. 49–65, 1999.
- [33] D. M. Kullmann, A. W. Moreau, Y. Bakiri, and E. Nicholson, “Plasticity of inhibition,” *Neuron*, vol. 75, no. 6, p. 951–962, 2012.

Supplementary Information

S1 Sampler derivation

The sampler was derived by combining an HMC step, and a Langevin step to add noise and ensure ergodicity. The most general equations describing HMC are given by

$$\tau_H \dot{\mathbf{u}} = \frac{\partial \log P(\mathbf{u}, \mathbf{v} | \mathbf{x})}{\partial \mathbf{v}} \quad (\text{S15a})$$

$$\tau_H \dot{\mathbf{v}} = -\frac{\partial \log P(\mathbf{u}, \mathbf{v} | \mathbf{x})}{\partial \mathbf{u}}. \quad (\text{S15b})$$

For the HMC, there is freedom to specify the distribution of the auxilliary variable, as we describe in the main text, we chose,

$$P(\mathbf{v} | \mathbf{u}, \mathbf{x}) = P(\mathbf{v} | \mathbf{u}) = \mathcal{N}(\mathbf{v}; \mathbf{u}, \mathbf{M}), \quad (\text{S16})$$

Substituting Equation (S16), into Equation (S15) gives,

$$\tau_H \dot{\mathbf{u}} = \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) \quad (\text{S17a})$$

$$\tau_H \dot{\mathbf{v}} = \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) - \frac{\partial \log P(\mathbf{u})}{\partial \mathbf{u}}. \quad (\text{S17b})$$

Next, we perform a Langevin step, by adding noise and going up the gradient of the log-probability,

$$\dot{\mathbf{u}} = \frac{1}{\tau_H \sigma_{\mathbf{v}}^2} \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) + \frac{1}{\tau_L} \frac{\partial \log P(\mathbf{u}, \mathbf{v} | \mathbf{x})}{\partial \mathbf{u}} + \sqrt{\frac{2}{\tau_L}} \boldsymbol{\eta}_{\mathbf{u}} \quad (\text{S18a})$$

$$\dot{\mathbf{v}} = \frac{1}{\tau_H \sigma_{\mathbf{v}}^2} \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) - \frac{1}{\tau_H} \frac{\partial \log P(\mathbf{u} | \mathbf{x})}{\partial \mathbf{u}} + \frac{1}{\tau_L} \frac{\partial \log P(\mathbf{u}, \mathbf{v} | \mathbf{x})}{\partial \mathbf{v}} + \sqrt{\frac{2}{\tau_L}} \boldsymbol{\eta}_{\mathbf{v}}. \quad (\text{S18b})$$

Substituting again for $P(\mathbf{v} | \mathbf{u})$ gives,

$$\dot{\mathbf{u}} = \frac{1}{\sigma_{\mathbf{v}}^2} \left(\frac{1}{\tau_H} - \frac{1}{\tau_L} \right) \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) + \frac{1}{\tau_L} \frac{\partial \log P(\mathbf{u} | \mathbf{x})}{\partial \mathbf{u}} + \frac{1}{\sqrt{2\tau_L}} \boldsymbol{\eta}_{\mathbf{u}} \quad (\text{S19a})$$

$$\dot{\mathbf{v}} = \frac{1}{\sigma_{\mathbf{v}}^2} \left(\frac{1}{\tau_H} + \frac{1}{\tau_L} \right) \mathbf{M}^{-1}(\mathbf{u} - \mathbf{v}) - \frac{1}{\tau_H} \frac{\partial \log P(\mathbf{u} | \mathbf{x})}{\partial \mathbf{u}} + \frac{1}{\sqrt{2\tau_L}} \boldsymbol{\eta}_{\mathbf{v}}. \quad (\text{S19b})$$

To obtain Equation (3), simply use $P(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{M})$, and $\mathbf{M} = \sigma_v^2 \mathbf{I}$.

Finally, Equation (3) requires the gradient of the likelihoods for the GSM, which are given by,

$$\frac{\partial \log P(\mathbf{x}|\mathbf{u}, z)}{\partial \mathbf{u}} = \frac{1}{\sigma_x^2} |z| \mathbf{A}^T (\mathbf{x} - |z| \mathbf{A} \mathbf{u}) \quad (\text{S20a})$$

and,

$$\frac{\partial \log P(\mathbf{x}|\mathbf{u}, z)}{\partial z} = \text{sign}(z) \mathbf{u}^T \mathbf{A}^T (\mathbf{x} - |z| \mathbf{A} \mathbf{u}) \quad (\text{S20b})$$

S2 Equivalence of filter and image models

We have two GSMs, labelled p for “pixel” and f for “filter”, so \mathbf{x}_p is the raw pixel intensities of the presented image, and $\mathbf{x}_f = \mathbf{F} \mathbf{x}_p$ is the filter activations. We therefore equate the variance of the posteriors (Equation (6)), which gives,

$$\mathbf{A}_p^T \mathbf{A}_p = \mathbf{A}_f^T \mathbf{A}_f. \quad (\text{S21})$$

If we rewrite, \mathbf{A}_p and \mathbf{A}_f in terms of left and right eigenvectors,

$$\mathbf{A}_p = \sum_i \lambda_{p,i} \mathbf{u}_{p,i} \mathbf{v}_{p,i}^T \quad (\text{S22})$$

$$\mathbf{A}_f = \sum_i \lambda_{f,i} \mathbf{u}_{f,i} \mathbf{v}_{f,i}^T, \quad (\text{S23})$$

where \mathbf{u} and \mathbf{v} are orthonormal, and the sum excludes any singular values that are 0, then Equation (S21) tells us that the eigenvalues and right eigenvectors of \mathbf{A}_p and \mathbf{A}_f are equal, assuming that the eigenvectors are all different. Without loss of generality, we can therefore choose our representation for \mathbf{A}_f to match our representation for \mathbf{A}_p ,

$$\lambda_{f,i} = \lambda_{p,i} = \lambda_i \quad (\text{S24})$$

$$\mathbf{v}_{f,i} = \mathbf{v}_{p,i} = \mathbf{v}_i \quad (\text{S25})$$

and the orthonormal basis $\mathbf{u}_{f,i}$ can be chosen arbitrarily. Equating the mean of the posteriors gives,

$$\mathbf{A}_p^T \mathbf{x}_p = \mathbf{A}_f^T \mathbf{x}_f = \mathbf{A}_f^T \mathbf{F} \mathbf{x}_p, \quad (\text{S26})$$

$$(\text{S27})$$

again, rewriting this in terms of the eigendecomposition gives,

$$\sum_i \lambda_i \mathbf{v}_i \mathbf{u}_{p,i}^T \mathbf{x}_p = \sum_i \lambda_i \mathbf{v}_i \mathbf{u}_{f,i}^T \mathbf{F} \mathbf{x}_p \quad (\text{S28})$$

It is sufficient to have,

$$\mathbf{u}_{p,i} = \mathbf{F}^T \mathbf{u}_{f,i} \quad (\text{S29})$$

rewriting this in terms of the matrices \mathbf{U}_f , and \mathbf{U}_p whose columns are $\mathbf{u}_{f,i}$ and $\mathbf{u}_{p,i}$, respectively,

$$\mathbf{F} = \mathbf{U}_f \mathbf{U}_p^T. \quad (\text{S30})$$

These results were checked by explicitly computing \mathbf{F} and \mathbf{A}_f that satisfied Equations (S21) and (S27) from \mathbf{A}_p .

S3 Robustness to the Langevin Timescale

The results in the main text are similar to those found with $\tau_L = 60$ ms and $\tau_L = 400$ ms, but quite different to those found with $\tau_L = 30$ ms and $\tau_L = 1000$ ms (Figure S1).

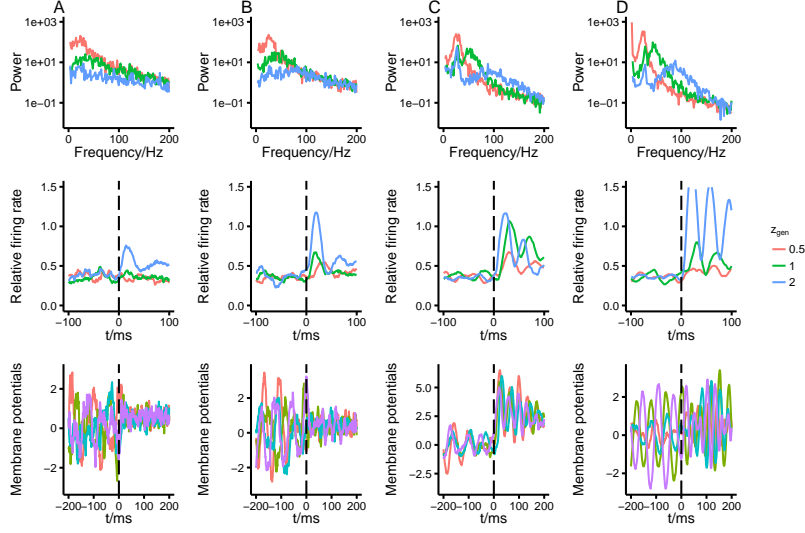


Figure S1: Our main results are robust to a range of τ_L . The top row is a power spectrum, the middle row displays the firing rate transient at stimulus onset, and the bottom row displays the membrane potential at stimulus onset for multiple trials and one neuron. The different lines in the first two rows correspond to different values of z_{gen} . In the bottom row, different lines correspond to different trials. **A** For $\tau_L = 30$ ms, transients are small or non-existent, and no clear trends are present in the peak frequency. **BC** For $\tau_L = 60$ ms (**B**), and $\tau_L = 400$ ms (**C**) the results are similar to those in the main text. **D** For $\tau_L = 1000$ ms, the results are quite different to those in the main text. In particular, the transient at stimulus onset lasts a long time, certainly longer than the observed value of around 50 ms.